Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view Posted by Herbert Poetzl on Fri, 30 Jun 2006 03:35:55 GMT

View Forum Message <> Reply to Message

```
On Thu, Jun 29, 2006 at 08:15:52PM -0400, jamal wrote:
> On Fri, 2006-30-06 at 09:07 +1200, Sam Vilain wrote:
> > jamal wrote:
>>> Makes sense for the host side to have naming convention tied
>>> to the guest. Example as a prefix: guest0-eth0. Would it not
>>> be interesting to have the host also manage these interfaces
>> via standard tools like ip or ifconfig etc? i.e if i admin up
>>> guest0-eth0, then the user in guest0 will see its eth0 going
> > > up.
> >
>> That particular convention only works if you have network namespaces
> > and UTS namespaces tightly bound.
> that would be one approach. Another less sophisticated approach is to
> have no binding whatsoever, rather some translation table to map two
> unrelated devices.
>> We plan to have them separate - so for
> > that to work, each network namespace could have an arbitrary
>> "prefix" that determines what the interface name will look like from
>> the outside when combined. We'd have to be careful about length
> > limits.
>> And guest0-eth0 doesn't necessarily make sense; it's not really an
> > ethernet interface, more like a tun or something.
> it wouldnt quiet fit as a tun device. More like a mirror side of the
> guest eth0 created on the host side
> i.e a sort of passthrough device with one side visible on the host (send
> from guest0-eth0 is received on eth0 in the guest and vice-versa).
> Note this is radically different from what i have heard Andrey and co
> talk about and i dont wanna disturb any shit because there seems to be
> some agreement. But if you address me i respond because it is very
> interesting a topic;->
thing is, we have several things we should care about
and some of them 'look' or 'sound' similar, although
they are not really ... I'll try to clarify
```

first, we want to have 'per guest' interfaces, which do not clash with any interfaces on the host or in

other guests

then, we want to 'connect' them, implicitely or explicetly with 'other' interfaces or devices inside other guests or on the host, here we have the following cases (some are a little special):

- lo interface, guest and host private (by default)
- tap/tun interfaces, again host/guest private
- tun like interfaces between host and guests
- tun like interfaces between guests
- 'normal' interfaces mapped into guests

on the traffic side we have the following cases:

- local traffic on the host
- local traffic on the guest
- local traffic between host and guest
- local traffic between guests
- routed traffic from guest via host
- bridged traffic from guest via host

special cases here would be tun/tap traffic inside a guest, but that can be considered local too

- >> So, an equally good convention might be to use sequential prefixes
- >> on the host, like "tun", "dummy", or a new prefix then a property
- > > of that is what the name of the interface is perceived to be to
- > > those who are in the corresponding network namespace.
- > >
- >> Then the pragmatic question becomes how to correlate what you see
- > > from `ip addr list' to guests.
- >
- > on the host ip addr and the one seen on the guest side are the same.
- > Except one is seen (on the host) on guest0-eth0 and another is seen
- > on eth0 (on guest).

this depends on the way the interfaces are handled and how they actually work, means:

if the interfaces _solely_ work via routing or bridging, then the 'host' end has to exist and be visible similar to 'normal' interfaces

if the traffic is (magically) mapped from guest interfaces to real (outside) host interfaces, we might want the same view as the guest has (i.e. basically a 'copy' which is not real)

> Anyways, ignore what i am saying if it is disrupting the discussion. IMHO input is always welcome .. helps the folks to do better thinking:) > cheers, > jamal > > >